

REMARKS

Claims 1 and 5 stand rejected under 35 USC §112, first paragraph and claims 4, 5, 12, 13, 14, and 17 stand rejected under 35 USC §112, second paragraph.

Claims 1-3, 6, 7, 10 and 11-14 stand rejected under 35 USC §102(b) as being anticipated by Wieland, U.S. patent 6,643,748. Claims 4, 5, 17 and 18 stand rejected under 35 USC §103(a) as being unpatentable over Wieland, U.S. patent 6,643,748 in view of Stai et al., U.S. patent 6,401,128. Claims 8, 9, 16 and 17 stand rejected under 35 USC §103(a) as being unpatentable over Wieland, U.S. patent 6,643,748 in view of Panas et al., U.S. patent 6,473,857.

Enclosed is a proposed corrected FIG. 1 including a changed label to HBA hardware 119 to distinguish from the SAN 120. The specification and claims have been amended herein to conform to the correct label of HBA hardware 119 and accommodate the Examiner's objections to the drawings noted by the Examiner.

Claims 1-6, 11-13 and 17 have been amended to more clearly state the invention and to accommodate the Examiner's rejections under 35 USC §112. Claim 7 has been cancelled. Reconsideration and withdrawal of the rejections under 35 USC §112 and the Examiner's objections to the drawings is respectfully requested.

Reconsideration and allowance of each of the claims 1-6 and 8-18, as amended, is respectfully requested. Each of the pending claims 1-6 and 8-18, as amended, is patentable over the references of record.

Wieland, U.S. patent 6,643,748 discloses a system and method are described to programmatically manage access between one or more nodes and a

plurality of associated devices, such as shared storage units. Each node is programmed to include a data structure, which identifies whether an associated device is to be within the scope of the respective node. The data structure may include persistent and/or temporary lists. Each device may be programmatically masked relative to the node by dynamically modifying the data structure of the node, such as by employing a predetermined interface. FIGS. 2A-2C are functional block diagrams representing features of the operating system (e.g., driver stack) of a node 12 which may be programmed and/or configured to mask communication paths between the node and devices (not shown) operatively coupled to the node. Referring to FIG. 2A, the node 12 includes a host bus adapter 28, which, for example, may be a SCSI bus adapter. The host bus adapter 28 includes a controller 30 configured and/or programmed to manage communication paths or channels between the node 12 and devices operatively connected to the node through an associated bus. Specifically, the controller 30 creates a device object 32a, 32b, 32c (collectively referred to as "device objects 32") for devices connected to the bus. Each device object 32 is a functional device object that identifies a device including, for example, its location (e.g., a logical unit number) and operating characteristics of the device. The term device is intended to be broadly construed so as to, for example, correspond to a hardware device or a functional unit of the device, such as a logical unit of data storage associated with a storage device.

Stai et al., U.S. patent 6,401,128 discloses a system and method for sending frames between a public device and a private device comprise a phantom

device mapping, an address translation, a frame payload translation, and a CRC regeneration. The system and method assign a phantom AL_PA for the public device and establishes a phantom device mapping between the phantom AL_PA and the public device's Port_ID. With the phantom device mapping, the disclosed system directs all communication between the public device and the private loop device as if the communication were between a phantom device and the private device.

Specifically, the system and method comprise a public-to-private address translation in one direction and a private-to-public address translation in the other direction. During the public-to-private address translation process, the source address of the frame is converted to a phantom AL_PA. The public-to-private address translation uses a Port_ID to phantom AL_PA mapping table and finds an entry where the Port_ID matches the source address. The public-to-private address translation replaces the source address of the frame with the phantom AL_PA of the matched entry, and the destination address with the AL_PA only of the destination device. The private-to-public address translation replaces the destination address of the frame with the Port_ID of the matched entry, and the source address with the fabric assigned address of the private device. As set forth at column 6, starting at line 25: Every frame between public device 106A and private loop device 110E requires an address translation, either a public-to-private or a private-to-public address translation. In addition, if the frame content (called payload) contains any address information, it is also changed accordingly. Typically, these types of frames are Extended Link Services (ELS) in Fibre Channel, and they are either ELS request or response frames. FIG. 4 is a functional

block diagram of a preferred embodiment of an ELS request payload translation. An ELS request payload translation process is required if the source address and/or the destination address of a frame being transmitted between a public device and a private device are part of the payload. An ELS request payload translation may be performed during a public-to-private translation or during a private-to-public translation. During the ELS request payload translation process, the frame type and Extended Link Services command code are examined to determine if payload translation is required. These are examined using an ELS request payload table which stores information on the frame type, command codes, and the information of fields to be modified. If payload translation is required, one or more fields in the frame payload is translated according to the ELS request payload table. If an ELS request payload translation is performed, the information for that frame is stored in a request payload cross-reference table. The request payload cross-reference table can then be used during an ELS response payload translation as described in more detail below with reference to FIG. 5.

Panas et al., U.S. patent 6,473,857 discloses a method for centralized and managed loading of boot images into one or more processors that are part of a file server for a mass storage system. In a computer system having at least one first controller, at least one input output processor (IOP), a first bus and a second bus, the present invention includes the steps of detecting readiness of the IOP to load a boot image, identifying across the first bus a location where the boot image will be loaded and loading the boot image across the second bus. The first controller may determine which of a plurality of boot images should be loaded. The first controller and the IOP

may each have first and second processors, with communication between the first processors being across the first bus and boot images being accessed by the second processors across the second bus. On the IOP, the first processor may control power to the second processor and may monitor the status of the second processor, reporting across the first bus to the first controller's first processor regarding the status of the IOP's second processor. The boot image may be copied to memory local to the IOP's second processor or it may be made available across the second bus. The boot image supplied may be adapted to normal, diagnostic, crash dump or other purposes. The progress of IOP booting is tracked and monitored. As stated at column 4, lines 36-61: The connection options 130 are various methods of connecting servers and clients to the ISAN server 102A. The serial connections 140 support network management, modems for remote management, and uninterruptible power supply messages. The front panel connection 142 supports a management connection with the front panel display of the ISAN server 102A. The Ethernet connection 144 supports an Ethernet interface for management protocols and for data transfer. The network interface 146 is one of potentially many high speed interfaces on the server. In some embodiments, the network interface 146 is a fibre channel interface with drivers for a fibre channel arbitrated loop (FC-AL). The network interface 146 may also include drivers for SCSI-3 over the fibre channel medium using fibre channel protocol (FCP). The hardware interface 126 provides interface specific hardware components. For example, the network interface 146 has a network interface specific set of software modules to support configuration, diagnostics, performance monitoring, and health and status

monitoring. The operating system 124, the tables 116, and the interfaces 118-122 support the virtual device and storage routing functionality of the ISAN server 102A. These components of the ISAN server 102A route storage transactions among appropriate storage options 128 and the connection options 130 using configured sets of driver modules in the system.

The present invention, as recited in independent claims 1 and 11, as amended, provides a novel storage area network (SAN) management and configuration method and apparatus via enabling in-band communications that solves a problem of some existing SAN arrangements. A problem exists in some known storage area network arrangements, for example, in a serial storage architecture (SSA), device driver writers and host based adapter (HBA) vendors provide a complex set of micro code calls. A management program would then interrogate the HBA, using micro code calls specific to the particular HBA vendor and model, then interpret the results in a way that is specific to that particular HBA vendor and model. One problem with this arrangement is that an in-depth understanding is needed for every HBA model of every vendor, which in the case of Fibre Channel, is impractical. There are too many vendors and too many models to implement this approach. The present invention, as recited in independent claims 1 and 11, as amended, provides a pass through in said HBA device driver for passing communications to a device in the storage area network from said SAN management application, including at least one topology analysis command, as expressly recited in amended independent claim 1 and the HBA device driver including at least one pass through service for passing a plurality of commands to said device in

the storage area network; said commands including at least one topology analysis command, as expressly recited in amended independent claim 11. This feature is not shown, nor suggested in the prior art references of record including Wieland, Panas et al., and Stai et al.

Each of the independent claims 1 and 11, as amended, is patentable over all the references of record including Wieland, Panas et al., and Stai et al. Independent claim 1, as amended, recites a storage area network (SAN) management and configuration method via enabling in-band communications comprising the steps of: utilizing a SAN management application for communicating with a host bus adapter (HBA) device driver, and providing a pass through in said HBA device driver for passing communications to a device in the storage area network from said SAN management application, including at least one topology analysis command.

Independent claim 11, as amended, recites a storage area network (SAN) management and configuration apparatus via enabling in-band communications comprising: a storage area network (SAN) management application for communicating with at least one SAN-connected host system; said SAN-connected host system including a management application agent for communicating with a host bus adapter (HBA) device driver; said HBA device driver for communicating with a device in the storage area network; said HBA device driver including at least one pass through service for passing a plurality of commands to said device in the storage area network; said commands including at least one topology analysis command.

The standard for lack of novelty, that is, for ``anticipation," is one of strict

identity. To anticipate a claim for a patent, a single prior source must contain all its essential elements. To anticipate under section 102, a prior art reference must disclose all the elements of the claimed invention or their equivalents functioning in essentially the same way. The Wieland patent does not disclose all the elements of the claimed invention of independent claims 1 and 11, as amended.

More specifically, the Wieland patent does not disclose a storage area network (SAN) management and configuration method via enabling in-band communications comprising the step of providing a pass through in said HBA device driver for passing communications to a device in the storage area network from said SAN management application, including at least one topology analysis command. Thus, the Wieland patent does not anticipate independent claim 1.

Similarly, the Wieland patent does not disclose a storage area network (SAN) management and configuration apparatus via enabling in-band communications comprising a HBA device driver for communicating with a device in the storage area network; and said HBA device driver including at least one pass through service for passing a plurality of commands to said device in the storage area network; said commands including at least one topology analysis command. Thus, the Wieland patent does not anticipate independent claim 11.

The references of record including Wieland, Panas et al., and Stai et al. do not render obvious the claimed invention and each of the independent claims 1 and 11, as amended, is patentable. No objective teaching in the prior art or knowledge generally available to one of ordinary skill in the art suggests the claimed subject matter

of independent claims 1 and 11, as amended. The references of record including Wieland, Panas et al., and Stai et al. fail to suggest or provide any objective teaching of a storage area network (SAN) management and configuration method and apparatus via enabling in-band communications as taught and claimed by Applicants. Only applicants teach providing a pass through in said HBA device driver for passing communications to a device in the storage area network from said SAN management application, including at least one topology analysis command, as expressly recited in amended independent claim 1 and the HBA device driver including at least one pass through service for passing a plurality of commands to said device in the storage area network; said commands including at least one topology analysis command, as expressly recited in amended independent claim 11. Thus, each of the independent claims 1 and 11 is patentable.

Dependent claims 2-6, 8-10, and 12-18 respectively depend from patentable independent claims 1 and 11 and further define the invention. Thus, each of the dependent claims 2-6, 8-10, and 12-18 is patentable.

Applicants have reviewed all the art of record, and respectfully submit that the claimed invention is patentable over all the art of record, including the references not relied upon by the Examiner for the rejection of the pending claims.

It is believed that the present application is now in condition for allowance and allowance of each of the pending claims 1-6 and 8-18 is respectfully requested. Prompt and favorable reconsideration is respectfully requested.

If the Examiner upon considering this amendment should find that a

Serial No. 09/657,234

telephone interview would be helpful in expediting allowance of the present application, the Examiner is respectfully urged to call the applicants' attorney at the number listed below.

Respectfully submitted,

By:



Joan Pennington
Reg. No. 30,885
Telephone: (312) 670-0736



FIG. 1

